

Admission Control and Loss Management for an Application-Level Statistical Service

Geoffrey G. Xie

Department of Computer Science
Naval Postgraduate School
Monterey, CA 93943
xie@cs.nps.navy.mil

Simon S. Lam

Department of Computer Sciences
The University of Texas at Austin
Austin, Texas 78712-1188
lam@cs.utexas.edu

Abstract

We present an admission control framework and loss management techniques in support of a guaranteed statistical service. The service is characterized by (i) the loss rate of *application data units* (ADUs) bounded below a specified value, and (ii) ADU losses distributed evenly among flows subscribing to the service and uniformly over the duration of each flow. Specifically, a flow is modeled as a sequence of bursts, each of which is a sequence of packets that carry the bits of an ADU. The first packet of each burst *carries information on the ADU* (e.g., its bandwidth requirement). This traffic model enables admission control at the burst level as well as at the flow level. Such a two level admission control approach is very effective in bounding end-to-end ADU loss rates of flows while maintaining high channel utilization in the network. The traffic model also enables simple techniques that can be used at a network channel to distribute ADU losses evenly among flows subscribing to the same statistical service, and to protect high priority ADUs (e.g., I frames of MPEG applications).

1 Introduction

In all packet switching networks, packets have a maximum size (in number of bits). Most application data units (ADUs) are too large to be carried in a single packet and must be segmented for network delivery. To an application, the end-to-end delays and loss rate of its data units are more relevant performance measures than ones specified for individual packets. For example, consider an application that

sends a sequence of video pictures over an IP network. Each picture may be segmented into a sequence of IP packets. Clearly the loss rate of the pictures is much more important to the application than the loss rate of the packets.

However, packet delays and packet loss rate are the performance measures used in the Internet service model [21, 23] and the ATM service architecture [9, 11]. Specifically, three types of services have been identified according to the quality of service (QoS) requirements for packets: deterministic (guaranteed), statistical (controlled-load), and best-effort. To provide better support to applications, we believe that the service models should be extended to include QoS requirements for ADUs.

Best-effort service is well understood, and it is provided by the current Internet. There is also a fairly good understanding of deterministic services following the recent development of rate-based packet scheduling algorithms (e.g., [17, 19]). However, many important questions remain in characterizing statistical services. The lack of a good understanding of statistical services is manifested in three areas:

- The loss rate of packets has been exclusively used as the performance measure. As discussed above, the loss rate of ADUs is much more relevant to an application. In particular, because of segmentation, even a small rate of packet losses could result in a large number of ADU losses.
- The term "statistical" is not well defined. There is a general consensus that it implies that a small fraction

of packets could be lost.¹ However, the consequence to applications can be quite different depending on how and at what time scale the loss rate is measured. For example, if the loss rate is measured at a large time scale, say over intervals of several hours, a large number of consecutive packets could be lost for a flow even when the measured loss rate is small. Such a statistical service (like the predictive service defined in [5]) is appropriate only for applications that can adapt to long periods of persistent losses. Moreover, the important issue of fair distribution of losses among flows has not received much attention.

- Complex statistical traffic models are often used in the development of new techniques to provide statistical performance guarantees at a small time scale (e.g., over intervals of milliseconds). Such models may not be practical because (i) real world traffic has long range dependency [10], and (ii) traffic monitoring and policing, if required, would be difficult to implement with a complex statistical traffic model.

This paper is concerned with providing a *guaranteed statistical service* to delay-sensitive flows. The term *statistical* means that the loss rate of the service may be nonzero. The loss rate is measured in application data units lost (including those late in arriving to their destinations). The term *guaranteed* means that the service's loss rate can be accurately predicted at connection setup time. In particular, we present an admission control framework and loss management techniques for providing the guaranteed statistical service.

We model a flow as a sequence of bursts, each of which is a sequence of packets that carry the bits of an ADU. The first packet of each burst carries information on the ADU (e.g., its bandwidth requirement). This traffic model enables admission control by a network node at the burst level as well as at the flow level. To achieve high channel utilization, overbooking is allowed in flow level admission control. Burst level admission control, on the other hand, is used to ensure that the capacity of each channel is not exceeded by

¹For a real-time service, packets that arrive to their destinations late (i.e., exceed their end-to-end delay bounds) are considered to be lost.

bandwidths allocated to flows such that delay guarantees can be provided to ADUs. As demonstrated by experimental results, such a two level admission control approach is very effective in bounding end-to-end ADU loss rates of flows while maintaining high channel utilization in the network. The traffic model also enables simple techniques that can be used at a channel to distribute losses evenly among flows subscribing to the same statistical service, and to protect high priority ADUs (e.g., I frames of MPEG applications).

A guaranteed statistical service, such as the one proposed herein, will greatly enhance the ability of existing networks to support interactive video and audio applications such as remote teaching and video conferencing, which cannot tolerate long durations of ADU (e.g. video picture or audio clip) losses, and for which a deterministic service is unnecessary and too costly.

The balance of this paper is organized as follows. In Section 2, our traffic model is described. In Section 3, an algorithm for single node admission control, at both the flow and burst levels, is presented. In Section 4, techniques for managing losses at a single node are presented. In Section 5, the issue of nodal allocation of an end-to-end ADU loss requirement is investigated. In Section 6, some related work is discussed.

2 Traffic Model

In designing networks that provide delay guarantees to variable bit rate (VBR) flows, we introduced the concept of a burst [16], which models a sequence of packets that encapsulate an application data unit. For video, for example, a burst models a sequence of packets that carry the encoded bits of a picture. The first and last packets of each burst are marked and the first packet carries information on the bit (or packet) rate of the burst.² In summary, we developed a *deterministic* flow model in which the boundaries and the bandwidth requirement of each ADU are encoded and accessible by the network.

²This part of our model can have different implementations. For example in ATM networks, it may be desirable to use RM cells to mark bursts and carry rate information.

Such a deterministic traffic model makes *burst-based rate allocation* possible. Specifically, at each channel, a reserved rate is not allocated to an admitted flow until the first packet of a burst arrives, and the rate is subsequently deallocated when the last packet of the burst departs [16]. (At any time, at most one burst in each flow is allocated its reserved rate.) As a result, the rate allocated to a VBR flow is variable, i.e., it changes from one burst to the next. Burst-based rate allocation has the advantage of allocating a rate that is *exactly* what a flow needs at all time.

The burst concept and traffic model are particularly applicable to network design in support of *application-level services* [16, 25]. Specifically, they allow each ADU's packets to be processed as a whole, and independently from other ADUs. Consequently much more predictable performance for ADUs can be rendered. Moreover, networks can be designed to provide ADUs with delay guarantees so long as the capacity of each channel is not exceeded by the aggregate rate allocated [16].

In what follows, we assume that each flow that subscribes to the guaranteed statistical service conforms to the deterministic traffic model defined below.

Definition 1 (deterministic traffic model) *A flow is modeled as a sequence of bursts, each of which is a sequence of packets that carry the bits of an ADU. The first and last packets of each burst are marked. The first packet carries some information on the ADU (including its bandwidth requirement).*

There are costs associated with requiring flows to conform to the above traffic model, namely: the extra bits for encoding ADU information in some packets, and the extra processing by a network node to retrieve and act on the information. However, they are bearable with today's network technologies (especially hardware), and the potential benefits outweigh the costs. In fact, my approach is consistent with the technology trend: It is advocated that the network be more active, i.e. to perform more application specific computations based on packet contents [8], so that better and more flexible services can be provided.

3 Admission Control by a Node

In this section, we describe an admission control algorithm that can be used by a node to bound the ADU loss rate at its channel below a specified value. We will discuss how to perform nodal allocation of an end-to-end loss requirement in Section 5.

For designing admission control, we assume that the path of each flow is fixed. For a particular flow, its path is a sequence of nodes, each of which consists of an outgoing channel, and a set of buffers for the channel where packets of different flows are queued.

Consider a particular node. Let C (bits/second) be the channel capacity dedicated to a guaranteed statistical service with a target ADU loss rate of p . (Note that it is straightforward to extend our design to a channel that is hierarchically shared by different agencies, by multiple classes of statistical services, etc. [1, 25]).

In addition to bounding the ADU loss rate, another goal of admission control is to achieve high channel utilization. Next, we present a novel approach that combines admission control at the flow level and admission control at the burst level to achieve both goals.

3.1 Burst Level Admission Control

Our traffic model, as discussed earlier, makes it possible for the packets of a burst (ADU) to be processed as a whole, and independently from other bursts. Thus, separate admission control can be performed for each burst. Specifically, a burst is accepted only if the burst's reserved rate³ does not exceed the channel's unallocated capacity; otherwise, the entire burst will be discarded. (Note that a similar mechanism, called the ATM Block Transfer (ABT) capability, is being standardized by ITU-T [14].)

With burst level admission control, packet losses are concentrated over a small number of ADUs, and channel capacity is not wasted on delivery of partial ADUs.⁴

Notation. In what follows, we use i to index flows, and use m to index bursts.

³The exact value is carried in the first packet of the burst.

⁴The usefulness of partial ADUs will be discussed later in Section 6.

Algorithm specification of burst level admission control

The variable A is used to store the aggregate rate allocated, and is initialized to 0. Let λ_m denote the bandwidth requirement (or reserved rate) of burst m .

- upon arrival of first packet of burst m

Burst_Admission_Control (m)

```

1      if (  $A + \lambda_m > C$  )
2          then    discard burst  $m$ ;
3          else    admit burst  $m$ ;
4                   $A := A + \lambda_m$ ;

```

- upon departure of last packet of burst m

$A := A - \lambda_m$;

The above algorithm has extremely low processing cost. Furthermore, it is performed only once per burst. Note that the average inter-burst arrival time is usually much larger than the average inter-packet arrival time. (This is especially true in an ATM network.) Therefore, the algorithm is suitable for high speed networks.

3.2 Flow Level Admission Control

Burst level admission control insures that channel capacity is not exceeded at any time by discarding bursts if necessary. This ensures that delay guarantees can be provided to bursts that are accepted by the node [16]. Since all ADU losses are due to burst level admission control, the goal of flow level admission control at the node becomes very specific: to allow as much overbooking as possible while bounding the probability that the channel's unallocated capacity is not sufficient for a newly arrived burst — denoted as the overflow probability — by p . The overflow probability and statistical multiplexing gains are closely related. In particular, if two channels have the same overflow probability, the utilization is higher, because of statistical multiplexing gains, for the one that has a larger capacity and is shared by more flows. Therefore, our flow level admission control exploits statistical multiplexing gains in an explicit manner.

Flow level admission control condition

Assume that the statistical service of the node is currently shared by a set of M flows (indexed by $1, 2, \dots, M$). Each flow, say i , supplies the following set of traffic (TSpec) parameters when making a reservation with the node: sustained bit (or packet) rate SR_i , peak rate PR_i , and rate variance RV_i . At any time, with burst-based rate allocation, at most one burst in each flow has its reserved rate allocated. Denote $\lambda_i(t)$ the reserved rate for the burst of flow i that is either allocated a rate or being processed by burst level admission control at time t . ($\lambda_i(t) = 0$ if there is no such burst.)

In [24], we derived, using a generalized version of central limit theorem, the following sufficient condition to bound the burst loss rate at the node below p :

$$\frac{Z_p}{Z} \leq 1, \quad (1)$$

where Z_p is the $(1 - p)$ percentile of the standard normal distribution, and

$$Z = \frac{C - \sum_{i=1}^M E[\lambda_i(t)]}{\sqrt{\sum_{i=1}^M Var[\lambda_i(t)]}}. \quad (2)$$

We refer to the value of (Z_p/Z) as the Statistical Multiplexing Intensity (SMI) of the service. It should never exceed the threshold of 1 to bound the burst loss rate below p . In practice, it is difficult to obtain the exact value of Z . However, Z can be estimated as follows:

$$\hat{Z} = \frac{C - \sum_{i=1}^M SR_i}{\sqrt{\sum_{i=1}^M RV_i}}. \quad (3)$$

In summary, our admission control algorithm accepts a new flow only if the following condition is not violated:

$$\frac{Z_p}{\hat{Z}} \leq 1. \quad (4)$$

Note that the source of flow i may not have a good estimate of RV_i at the time of connection setup. In such a case, RV_i is upper bounded by $SR_i \cdot (PR_i - SR_i)$, which can be used as a pessimistic estimate.⁵

⁵On-line measurement of SR and RV can be performed for admitted flows to improve the performance of flow level admission control. We are currently investigating such techniques.

Algorithm specification of flow level admission control

The variables B and V are used to store respectively the unallocated channel capacity and the aggregate rate variance of admitted flows. Initially B is set to C , and V is set to 0. We assume that if a source does not have a good estimate of RV at the time of connection setup, it will let the network know by setting RV to -1 .

- upon receiving connection request from flow i

Flow_Admission_Control (i)

```

1  if ( $B - SR_i \leq 0$ )
4    then reject the flow;

5  if ( $RV_i = -1$ )
6    then  $RV_i := SR_i * (PR_i - SR_i)$ ;
7   $SMI := Z_p * (\text{sqrt}(V + RV_i) / (B - SR_i))$ ;
8  if ( $SMI \leq 1$ )
9    then accept the flow;
10      $B := B - SR_i$ ;
11      $V := V + RV_i$ ;
12  else reject the flow;
```

3.3 Experimental Results




Figure 1: Admission control for single node

We have evaluated the single node admission control algorithm by performing a set of simulation experiments. (More details about the experiments can be found in [24].) The simulation configuration is shown in Figure 1. The nodes labeled by VS denote video sources, and VD their destination. Each video source generates 53-byte packets from a trace file obtained from a MPEG video sequence, and each MPEG frame (or picture) is considered an ADU.

The two level admission control algorithm is implemented for channel L1 with a target ADU loss rate of p . Each video source makes a reservation with L1, and starts sending packets to the network only after the reservation is successful. Packets are scheduled based on their virtual clock values [27]. The channel capacity of L1 as well as the value of p were varied in different experiments. We ran each experiment for 10 seconds of simulated time.

Channel utilization

In Figure 2, we plot the channel utilization as a function of the target picture loss rate p . The result shows that the channel is used much more efficiently with a statistical service than a deterministic service (with zero loss rate). The price to pay is a small non-zero picture loss probability. The utilization increase is more significant with a higher channel capacity, from below 30% to above 70% in the case where the capacity of L1 is 56 Mbps. This is because the improvement is due to statistical multiplexing gains, which are larger with more flows sharing the channel.




Figure 2: Channel utilization improvement

Actual picture loss rate

For the channel utilization gain to be meaningful, the actual ADU (picture) loss rates in the experiments cannot be much larger than their respective target values. In Figure 3, we compare the actual picture loss rate in each experiment, averaged over five simulation runs using different random seeds, with the target value. From the figure, we conclude




Figure 3: Actual vs. target picture loss rate

that our admission control algorithm predicts the actual loss rate well when a large number of flows share the channel. (Around 30 flows were admitted when the channel capacity of L1 was set to 56 Mbps.) This agrees with our analysis; the larger the number of flows sharing the channel, the better approximation based on the central limit theorem. Note that the solid 45 degree line represents perfect prediction by the central limit theorem.

4 Loss Management at a Node

For the guaranteed statistical service to be viable, considering the fact the service will be shared by many flows, the issue of loss distribution must be addressed. In particular, the ADU losses should be distributed evenly among all flows subscribing to the service and uniformly over the duration of each flow.

In this section, we will discuss how to modify the admission control algorithm to enable the channel to anticipate and manage ADU losses.⁶

We want to achieve two goals with such *active loss management*. The first goal is clear from the requirement of a guaranteed statistical service; that is, losses should be distributed as evenly as possible among flows that share the channel. The second goal is motivated by the fact that to many applications, some of their ADUs are more important than others. For example, to an MPEG decoder, I frames are more important than either P or B frames. With our traffic

⁶We do not consider packet losses due to link or transmission errors because the probability of their occurrence is extremely small with today's hardware, and they can be dealt with by link layer protocols.

model, applications can easily mark important ADUs and request that the network give priority to them. Therefore, it is desirable that loss management at a node can facilitate protection of high priority ADUs marked by applications.

Next, we present two loss management techniques based on selective early ADU discard (SEAD). Similar to early packet discard proposals [20, 22], SEAD achieves the goals of fair loss distribution and protection of high priority ADUs by taking early control actions. Specifically, a trigger point is set at αC , where $0 < \alpha < 1$, and control actions are triggered whenever the aggregate rate allocated by the channel exceeds αC .

4.1 SEAD-1

For the first technique (named SEAD-1), the flow level admission control algorithm remains the same as the one described in Section 3.2. The burst level admission control algorithm is modified to implement SEAD. In particular, there are two control actions. The first involves discarding unmarked ADUs to protect high priority ADUs, and the second makes use of binary counters, one for each flow (c_i for flow i), to prevent a flow from losing consecutive ADUs. Below is a specification of the modified burst level admission control algorithm. Initially the binary counter of each flow is set to zero. $f(m)$ denotes the index of the flow to which burst m belongs.

Burst_Admission_Control (α, m)

```

1  if (  $A + \lambda_m > C$  )
2      then discard burst  $m$ ;
3      else if (  $A + \lambda_m > \alpha C$  and
                burst  $m$  not marked with high priority )
4          then if (  $c_{f(m)} = 0$  )
5              then discard burst  $m$ ;
6                   $c_{f(m)} := 1$ ;
7              else admit burst  $m$ ;
8                   $A := A + \lambda_m$ ;
9                   $c_{f(m)} := 0$ ;
10     else admit burst  $m$ ;
11          $A := A + \lambda_m$ ;

```

Note that in SEAD-1, flow level admission control is not modified. The advantage is that high utilization can be maintained at the same time that high priority ADUs have a loss rate much smaller than p . This advantage does come with a price. Namely, those unmarked ADUs can experience more losses. However, as stated in the theorem below, the overall ADU loss rate is still bounded, albeit by a value a little larger than p .

Theorem 1 *Let β denote the maximum channel utilization achievable by flow level admission control. When SEAD-1 is used, the overall ADU loss rate is bounded by p_1 where p_1 is the value such that the $(1 - p_1)$ percentile of the standard normal distribution equals $\frac{\alpha - \beta}{1 - \beta} Z_p$.*

The proof is omitted due to space limit. It can be found in [24].

Example 1 *Consider the case where $\alpha = 0.95$, $\beta = 0.7$, and $p = 0.2\%$. We have*

$$\frac{\alpha - \beta}{1 - \beta} Z_p = \frac{0.95 - 0.7}{1 - 0.7} \cdot 2.87 = 2.4$$

Therefore, the overall loss rate is still bounded by 0.8% when SEAD-1 is used.

We have evaluated the performance of SEAD-1 by simulation experiments. The simulation configuration shown in Figure 1 was used. The channel capacity of L1 is 48 Mbps. For the experiments, I frames are marked high priority, and all other frames are unmarked. A pair of experiments, labeled BASE and SEAD-1, were performed for every simulation run, and in each run a different random seed was used. The original admission control algorithm (i.e. without using SEAD) is used by L1 in BASE experiments. Table 1 contains the result from a typical simulation run. The result shows that by using SEAD-1, the amount of high priority ADU losses is reduced by more than 50%, and the losses are distributed more evenly among flows (as indicated by the fact that the highest loss rate for one flow is much smaller).

SEAD does add some processing cost to the admission control algorithm. However, since the extra processing is performed on a per ADU basis, the associated cost is not

	Channel Util	ADU losses		Highest Loss Rate for One Flow
		Total	High Pri	
BASE	69 %	16	8	2.1%
SEAD-1	69 %	54	3	1.36%

Table 1: L1 capacity = 48 Mbps, $p = 0.2\%$, $\alpha = 0.95$

significant because the average ADU size is quite large compared to the average packet size. Note that because of burst level admission control, the units for statistical multiplexing become ADUs instead of packets. That explains why simple binary counters are effective in spreading out ADU losses among flows as we have observed in the experiments.

4.2 SEAD-2

The second technique (named SEAD-2) is almost the same as SEAD-1 except that it requires a small but significant modification to the flow level admission control algorithm. Specifically, αC is used in place of C in flow level admission control, i.e., B is initialized to αC instead of C in the algorithm specification given in Section 3.2.

Compared to SEAD-1, the maximum channel utilization achievable with SEAD-2 is lower. However, the difference can be made insignificant by choosing a α value close to 1, say 0.95. Our experimental results indicate that such a α value is sufficient for reaping all the benefits of SEAD-2, namely:

- loss rate of high priority ADUs significantly reduced,
- loss rate of all ADUs below target value, and
- losses distributed more evenly among flows.

We have evaluated SEAD-2 using simulation experiments. Table 2 contains the result of a typical simulation run. The result shows that while the achieved channel utilization is a little lower, both the amount of high priority ADU losses and the highest loss rate for a flow are significantly lower when SEAD-2 is used.

5 Nodal Allocation of End-to-End Loss Requirement

So far we have shown that a network node can use admission control and loss management to provide application-level statistical services. What an application really needs,

	Channel Util	ADU losses		Highest Loss Rate for One Flow
		Total	High Pri	
BASE	71%	34	17	3.1%
SEAD-1	71%	86	10	2.7%
SEAD-2	67.6%	28	4	1.03%

Table 2: L1 capacity = 48 Mbps, $p = 0.5\%$, $\alpha = 0.95$

however, is an end-to-end loss performance guarantee. In this section, we discuss how to extend our design from a single node to a flow path. In particular, we address the problem of nodal allocation of an end-to-end ADU loss requirement.

The nodal allocation problem can be formulated as follows: Consider a flow that requires a bound of F on its end-to-end ADU loss rate. Denote H the number of nodes in its flow path. (The source and destination are not considered part of the flow path.) Clearly each node h , $h = 1, 2, \dots, H$, must guarantee a bound, F_h , on the ADU loss rate of its channel. The question is how to choose these F_h values to bound the end-to-end loss rate below F while achieving good network performance.

Similar problems have been studied in the context of packet (or cell) loss performance. In [7], it is shown that

$$F \approx \sum_{h=1}^H F_h \quad (5)$$

when $F \ll 1$, and if cell losses at each node are independent. In [18], an argument has been made that for a stringent end-to-end loss constraint (i.e., $F \ll 1$), a non-uniform distribution of loss among nodes does not significantly improve network performance.

ADUs usually have much larger sizes than packets or cells. Therefore the results cited above may not be applicable to application-level service. In light of this uncertainty, we have performed a set of simulation experiments on path loss performance. The simulated network is shown in Figure 4. There are six video flows (V1-V6) that traverse through a path of 4 nodes (L1-L4) to their destination. The channel of each node has a capacity of 40 Mbps. The admission control algorithm specified in Section 3 is implemented at each node. There is also cross traffic carrying video data through

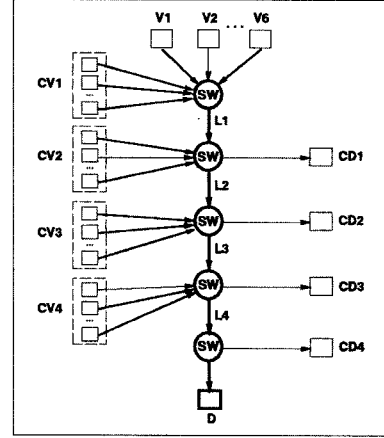


Figure 4: A simulated network path.

each node (e.g., CV1 flows to L1 to CD1). Real life MPEG video sequences are again used to generate traffic traces.

In particular, we would like to obtain results that help us answer the following two questions:

1. Is (5) still true in the context of ADU loss performance?
2. Is non-uniform distribution of loss among nodes useful?

5.1 Effect of Source Thinning

To answer the first question, we designed an experiment to study the effect of source thinning. Source thinning refers to the fact that a flow may lose some traffic intensity due to upstream losses as it traverses downstream. Specifically, we chose an identical target picture loss rate of 0.2% for each node, i.e., $p_1 = p_2 = p_3 = p_4 = 0.2\%$. The result is shown in Figure 5, where the percentage of total losses at each node is illustrated using a bar graph. Note that the losses are averaged over 20 simulation runs using different random seeds. The results show that source thinning does not have a significant impact on ADU losses of a path when there is a large amount of cross traffic through each node. In such a case (likely for future high speed networks), it is not overly conservative to assume that ADU losses of a flow are independent at each node. Therefore (5) is applicable to our framework of application-level service.




Figure 5: Losses at different nodes in a path.




Figure 6: Non-uniform nodal allocation.

5.2 Usefulness of Non-Uniform Nodal Allocation

To answer the second question, we repeated the previous experiment with one simple modification. Specifically, we assign different target loss rates to the four nodes: 0.3% to the two upstream nodes (L1 and L2) and 0.1% to the two downstream ones. As expected, the total losses, averaged over 20 simulation runs, are very close for these two experiments (46 losses in the new experiment vs. 48 in the original one). The percentage of total losses at each node for the new experiment is illustrated using a bar graph in Figure 6. The results indicate that given an end-to-end loss requirement, the ADU losses at downstream nodes can be reduced using non-uniform nodal allocation. Such reduction can be quite desirable for the purpose of minimizing wasteful work at the upstream nodes. (A similar observation has motivated proposals to give high priority to packets that have large hop counts.) Because ADUs usually have large sizes, the potential gain of non-uniform nodal allocation, measured by ADU

losses, can be significant. However as far as the overall network utilization is concerned, non-uniform nodal allocation does not cause any significant change.

6 Related Work

In this section, we review some related work in the area of application-level QoS. More discussion on related work can be found in [24].

Adaptive Resource Reservation

One main feature of the proposed deterministic traffic model is that it enables resource reservation on a per ADU basis for delay-sensitive VBR traffic. The need for such adaptive resource reservation was also independently discovered by several other researchers. Zhang and Knightly have proposed a renegotiated deterministic VBR service, called RED-VBR, based on their D-BIND traffic model [26]. Specifically, a compressed video source renegotiates its traffic specification and QoS with the network when its flow rate changes significantly. When the renegotiation fails, the source adjusts the quality control factor of the video encoder to lower the transmission rate. Grossglauser et. al. have designed a Renegotiated-Constant Bit Rate (R-CBR) service for VBR flows using a similar approach [12]. The main difference is that R-CBR builds the renegotiation service on top of a constant bit rate service. Chong et. al. have concentrated on the on-line prediction of the dynamic bandwidth requirement of a VBR flow using neural networks [3]. Additionally, several researchers (e.g., in [2]) have proposed to use in-call resource renegotiation as a fast reservation mechanism for bursty data traffic.

Usefulness of Partially Delivered ADUs

Our work has been motivated by the distinction between packets and ADUs. In particular, it is assumed that the entire ADU would be considered lost by the application if one of the ADU's packets is lost. It should be noted that there have been work such as the application level framing (ALF) proposal by Clark and Tennenhouse [6] that permits a partial ADU, i.e. one with some packet(s) lost, to be delivered to the application. While partially delivered ADUs are useful to some applications, they cannot be utilized by

delay-sensitive applications without some QoS degradation. It is our view that the network should have the capabilities of (1) meeting the requirements of applications that cannot tolerate QoS degradation, and (2) limiting the degradation as much as possible for those that can. That is why in this work, partially delivered ADUs are considered the same as losses. Similarly, late ADUs are considered lost even though there are applications (e.g., Mbone audio/video tools) that can adapt to change of delays.

References

- [1] Jon C.R. Bennett and Hui Zhang. Hierarchical packet fair queueing algorithms. In *Proceedings of ACM SIGCOMM '96*, pages 143–156, Stanford University, CA, August 1996.
- [2] E.E. Boyer and D.P. Tranchier. A reservation principle with applications to the ATM traffic. *Computer Networks and ISDN Systems*, 24:321–334, 1992.
- [3] Song Chong, San-Qi Li, and Joydeep Ghosh. Predictive dynamic bandwidth allocation for efficient support of real-time VBR video over ATM. *IEEE Journal on Selected Areas in Communications*, 13:12–23, January 1995.
- [4] I. Cidon and R. Guerin. An investigation of application level performance in ATM networks. In *Proceedings of IEEE INFOCOM '95*, pages 129–137, Boston, MA, March 1995.
- [5] David Clark, Scott Shenker, and Lixia Zhang. Supporting real-time applications in an integrated services packet network: Architecture and mechanism. In *Proceedings of ACM SIGCOMM '92*, pages 14–26, August 1992.
- [6] David Clark and David L. Tennenhouse. Architectural considerations for a new generation of protocols. In *Proceedings of ACM SIGCOMM '90*, pages 200–208, August 1990.
- [7] G. de Veciana, G. Kesidis, and J. Walrand. Resource management in wide-area ATM networks using effective bandwidth. *IEEE Journal on Selected Areas in Communications*, August 1995. *Special issue on Advances in Fundamentals of Networking*.
- [8] D.L. Tennenhouse *et al.* A survey of active network research. *IEEE Communications*, January 1997.
- [9] The ATM Forum. *ATM Traffic Management Specification*. 1995. Version 4.0.
- [10] M. Garrett and W. Willinger. Analysis, modeling and generation of self-similar VBR video traffic. In *Proceedings of ACM SIGCOMM '94*, pages 269–280, August 1994.
- [11] Mark W. Garrett. A service architecture for ATM: From applications to scheduling. *IEEE Network Magazine*, pages 6–14, May 1996.
- [12] M. Grossglauser, S. Keshav, and D. Tse. RCBP: A simple and efficient service for multiple time-scale traffic [extended version]. *An early version in ACM SIGCOMM '95*.
- [13] R. Guerin, H. Ahmadi, and M. Naghshineh. Equivalent capacity and its application to bandwidth allocation in high-speed networks. *IEEE Journal on Selected Areas in Communications*, 9(7):968–981, 1991.
- [14] ITU-T Rec. I.371. *Traffic Control and Congestion Control in B-ISDN*. Perth, U.K., Nov. 6–14, 1995.
- [15] Sugih Jamin, Peter Danzig, Scott Shenker, and Lixia Zhang. A measurement-based admission control algorithm for integrated services packet networks. In *Proceedings of ACM SIGCOMM '95*, pages 1–13, August 1995.
- [16] Simon S. Lam and Geoffrey G. Xie. Burst scheduling networks. *Performance Evaluation*, 1997. An abbreviated version in *Proceedings of INFOCOM '95*, April 1995.
- [17] Simon S. Lam and Geoffrey G. Xie. Group priority scheduling. *IEEE/ACM Transactions on Networking*, 5(2):205–218, April 1997. An abbreviated version in *Proceedings of INFOCOM '96*, March 1996.
- [18] R. Nagarajan. *Quality-of-Service Issues in High-speed Networks*. PhD thesis, University of Massachusetts at Amherst, August 1993.
- [19] Abhay K. Parekh and Robert G. Gallager. A generalized processor sharing approach to flow control in integrated services networks - the multiple node case. In *Proceedings of IEEE INFOCOM '93*, pages 521–530, March 1993.
- [20] A. Romanow and S. Floyd. The dynamics of TCP traffic over ATM networks. In *Proceedings of ACM SIGCOMM '94*, London, England, August 1994.
- [21] S. Shenker, C. Partridge, and R. Guerin. Specification of guaranteed quality of service. Technical report, February 1997. Internet Draft, *draft-ietf-intserv-guaranteed-svc-07.txt*.
- [22] J.S. Turner. Maintaining high throughput during overload in ATM networks. In *Proceedings of IEEE INFOCOM '96*, pages 287–295, San Francisco, CA, April 1996.
- [23] J. Wroclawski. Specification of the controlled-load network element service. Technical report, November 1996. Internet Draft, *draft-ietf-intserv-ctrl-load-svc-04.txt*.
- [24] Geoffrey G. Xie and Simon S. Lam. Admission control and loss management for an application-level statistical service. Technical Report CS-97-Xg-1, Naval Postgraduate School, January 1997. Available from <http://www.cs.nps.navy.mil/people/faculty/xie/pub>.
- [25] Geoffrey G. Xie and Simon S. Lam. Real-time block transfer under a link sharing hierarchy. In *Proceedings of IEEE INFOCOM '97*, Kobe, Japan, April 1997.
- [26] Hui Zhang and Edward W. Knightly. RED-VBR: A renegotiation-based approach to support delay-sensitive VBR video. To Appear in *ACM Multimedia Systems Journal*. An early version in *Proceedings of NOSSDAV '95*, Durham, NH, April 1995.
- [27] Lixia Zhang. VirtualClock: A new traffic control algorithm for packet switching networks. In *Proceedings of ACM SIGCOMM '90*, pages 19–29, August 1990.